

# Diversity Leads to Generalization in Neural Networks

Bo Xie\*, Yingyu Liang†, Le Song‡

## Abstract

Neural networks are a powerful class of functions that can be trained with simple gradient descent to achieve state-of-the-art performance on a variety of applications. Despite their practical success, there is a paucity of results that provide theoretical guarantees on why they are so effective. Lying in the center of the problem is the difficulty of analyzing the non-convex objective function with potentially numerous local minima and saddle points. Can neural networks corresponding to the stationary points of the objective function learn the true labeling function? If yes, what are the key factors contributing to such generalization ability?

In this paper, we provide answers to these questions by analyzing one-hidden-layer neural networks with ReLU activation, and show that despite the non-convexity, neural networks with diverse units can learn the true function. We bypass the non-convexity issue by directly analyzing the first order condition, and show that the loss is bounded if the smallest singular value of the “extended feature matrix” is large enough. We make novel use of techniques from kernel methods and geometric discrepancy, and identify a new relation linking the smallest singular value to the spectrum of a kernel function associated with the activation function and to the diversity of the units. Our results also suggest a novel regularization function to promote unit diversity for potentially better generalization ability.

## 1 Introduction

Neural networks are a powerful class of nonlinear functions which have been successfully deployed in a variety of machine learning tasks. In the simplest form, neural networks with one hidden layer are linear combinations of nonlinear basis functions (units),

$$f(x) = \sum_{k=1}^n v_k \sigma(w_k^\top x) \quad (1)$$

where  $\sigma(w_k^\top x)$  is a basis function with weights  $w_k$ , and  $v_k$  is the corresponding combination coefficient. Learning with neural networks involves adapting both the combination coefficients and the basis functions at the same time, usually by minimizing the empirical loss

$$L(f) = \frac{1}{2m} \sum_{l=1}^m \ell(y_l, f(x_l)) \quad (2)$$

with first-order methods such as (stochastic) gradient descent. It is believed that basis function adaptation is a crucial ingredient for neural networks to achieve more compact models and better performance [Barron, 1993, Yang et al., 2014].

However, the empirical loss minimization problem involved in neural network training is non-convex with potentially numerous local minima and saddle points. This makes formal analysis of training neural networks very challenging, and hence few theoretical results are available for their generalization ability. Given the empirical success of neural networks, a sequence of important and urgent scientific questions need to be investigated: Can neural networks corresponding to stationary points of the empirical loss learn the true labeling function? If the answer is yes, then what are the key factors contributing to their generalization ability? Based on these understandings, can we design better regularization schemes and learning algorithms to improve the training of neural networks?

In this paper, we provide partial answers to these questions by analyzing one-hidden-layer neural networks with rectified linear units (ReLU) in a least-squares regression setting. We show that despite the non-convexity, neural

\*College of Computing, Georgia Institute of Technology. Email: bo.xie@gatech.edu

†Department of Computer Science, Princeton University. Email: yingyu@cs.princeton.edu

‡College of Computing, Georgia Institute of Technology. Email: lsong@cc.gatech.edu

networks with diverse units can learn the true function. More specifically, we show that the training loss of neural networks decreases in proportion to  $\|\partial L/\partial W\|^2/s_m^2(D)$  where  $\partial L/\partial W$  is the gradient and  $s_m(D)$  is the minimum singular value of the extended feature matrix  $D$  (defined in Section 3.1). The minimum singular value is lower bounded by two terms, where the first term is related to the spectrum property of the kernel function associated with the activation  $\sigma(\cdot)$ , and the second term quantifies the diversity of the units, measured by the classical notion of geometric discrepancy of a set of vectors. Essentially, the slower the decay of the spectrum, the better the generalization ability; the more diverse the unit weights, the more likely stationary points will result in small training loss and generalization error.

We bypass the hurdle of non-convexity by directly analyzing the first order condition of the learning problem, which implies that the training loss is bounded if the minimum singular value of the extended feature matrix is large enough. Bounding the singular value is challenging because it entangles the nonlinear activation function, the weights and data in a complicated way. Unlike most previous attempts, we directly analyze the effect of nonlinearity without assuming independence of the activation patterns from actual data; in fact, the dependence of the patterns on the data and the unit weights underlies the key connection to activation kernel spectrum and the diversity of the units.

We have constructed a novel proof, which makes use of techniques from geometric discrepancy and kernel methods, and have identified a new relation linking the smallest singular value to the diversity of the units and the spectrum of a kernel function associated with the unit. More specifically,

- We identify and separate two factors in the minimum singular value: 1) an ideal spectrum that is related to the kernel of the activation function and an ideal configuration of diverse unit weights; 2) deviation from the ideal spectrum measured by how far away actual unit weights are from the diverse configuration. This new perspective reveals benign conditions in learning neural networks.
- We characterize the deviation from the ideal diverse weight configuration using the concept of discrepancy, which has been extensively studied in the geometric discrepancy theory. This reveals an interesting connection between the discrepancy of the weights and the training loss of neural networks. Therefore, it serves as a clean tool to analyze and verify the learning and the generalization ability of the networks.

Our results also suggest a novel regularization scheme to promote unit diversity for potentially better generalization. In [?], it is shown that diversity of the neurons leads to smaller network size and better performance.

Whenever possible, we corroborate our theoretical analysis with numerical simulations. These numerical results include computing and verifying the relationship between the discrepancy of a learned neural network and the minimum singular value. Additionally, we measure the effects on the discrepancy with and without regularization. In all these examples, the experiments match with the theory nicely and they accord with the practice of using gradient descent to learn neural networks.

## 2 Related work

Kernel methods have many commonalities with one-hidden-layer neural networks. The random feature perspective [Rahimi and Recht, 2009, Cho and Saul, 2009] views kernels as linear combinations of nonlinear basis functions, similar to neural networks. The difference between the two is that the weights are random in kernels while in neural networks they are learned. Using learned weights leads to considerable smaller models as shown in [Barron, 1993]. However it is a non-convex problem and it is difficult to find the global optima. *e.g.*, one-hidden-layer networks are NP-complete to learn in the worst case [Blum and Rivest, 1993]. We will make novel use of techniques from kernel methods to analyze learning in neural networks.

The empirical success of training neural networks with simple algorithms such as gradient descent has motivated researchers to explain their surprising effectiveness. In [Choromanska et al., 2015], the authors analyze the loss surface of a special random neural network through spin-glass theory and show that for many large-size networks, there is a band of exponentially many local optima, whose loss is small and close to that of a global optimum. The analyzed polynomial network is different from the actual neural network being used which typically contains ReLU nowadays. Moreover, the analysis does not lead to a generalization guarantee for the learned neural network.

A similar work shows that all local optima are also global optima in linear neural networks [Kawaguchi, 2016]. However their analysis for nonlinear neural networks hinges on independence of the activation patterns from the actual data, which is unrealistic. Some other works try to argue that gradient descent is not trapped in saddle points [Lee et al., 2016, Ge et al., 2015], as was suggested to be the major obstacle in optimization [Dauphin et al., 2014].

There is also a seminal work using tensor method to avoid the non-convex optimization problem in neural network [Janzamin et al., 2015]. However, the resulting algorithm is very different from typically used algorithms where only gradient information of the empirical loss  $L$  is used.

[Soudry and Carmon, 2016] is the closest to our work, which shows that zero gradient implies zero loss for all weights except an exception set of measure zero. However, this is insufficient to guarantee low training loss since small gradient can still lead to large loss. Furthermore, their analysis does not characterize the exception set and it is unclear a priori whether the set of local minima fall into the exception set.

### 3 Problem setting and preliminaries

We will focus on a special class of data distributions where the input  $x \in \mathbb{R}^d$  is drawn uniformly from the unit sphere,<sup>1</sup> and assume that  $|y| \leq Y$ . We consider the following hypothesis class:

$$\mathcal{F} = \left\{ \sum_{k=1}^n v_k \sigma(w_k^\top x) : v_k \in \{\pm 1\}, \sum_{k=1}^n \|w_k\| \leq C_W \right\} \quad (3)$$

where  $\sigma(\cdot) = \max\{0, \cdot\}$  is the rectified linear unit (ReLU) activation function,  $\{w_k\}$  and  $\{v_k\}$  are the unit weights and combination coefficients respectively,  $n$  is the number of units, and  $C_W$  is some constant. We restrict  $v_k \in \{-1, 1\}$  due to the positive homogeneity of ReLU,

$$f(x) = \sum_{k=1}^n v_k \sigma(w_k^\top x) = \sum_{k=1}^n \frac{v_k}{|v_k|} \sigma(|v_k| w_k^\top x).$$

That is, the magnitude of  $v_k$  can always be scaled into the corresponding  $w_k$ . For convenience, let

$$W := (w_1^\top, w_2^\top, \dots, w_n^\top)^\top \quad (4)$$

be the column concatenation of the unit parameters; also let  $W$  denote the set  $\{w_i : 1 \leq i \leq k\}$ . Let

$$\mathcal{F}_W = \left\{ W : \sum_k \|w_k\| \leq C_W \right\} \quad (5)$$

denote the feasible set of  $W$ 's. A function  $f \in \mathcal{F}$  will depend on  $v$  and  $W$ , and it can be written as  $f(x; v, W)$ . But when clear from the context, it is shorten as  $f(x)$ .

Given a set of *i.i.d.* training examples  $\{x_l, y_l\}_{l=1}^m \subseteq \mathbb{R}^d \times \mathbb{R}$ , we want to find a function  $f \in \mathcal{F}$  which minimizes the following least-squares loss function<sup>2</sup>

$$L(f) = \frac{1}{2m} \sum_{l=1}^m (y_l - f(x_l))^2.$$

Typically, gradient descent over  $L(f)$  is used to learn all the parameters in  $f$ , and a solution with small gradient is returned at the end.<sup>3</sup> However, adjusting the bases  $\{w_k\}$  leads to a non-convex optimization problem, and there is no theoretical guarantee on the training and test performance.

Our primary goal is to identify conditions for bounding the training loss and generalization error. We need to identify a set  $\mathcal{G}_W \subseteq \mathcal{F}_W$  such that when gradient descent outputs a solution  $W \in \mathcal{G}_W$  with the gradient norm  $\|\partial L / \partial W\|$  smaller than  $\epsilon$ , then the training and test errors can be bounded by  $O(\epsilon^2)$ . Ideally,  $\mathcal{G}_W$  should have clear characterization that can be easily verified, and should contain most  $W$  in the parameter space (especially those solutions obtained in practice).

On notation, we will use  $c, c'$  or  $C, C'$  to denote constants and its value may change from line to line.

<sup>1</sup>The restriction of input to the unit sphere is not too stringent since many inputs are already normalized. Furthermore, it is possible to relax the uniform distribution to sub-gaussian rotationally invariant distributions, but we use the current assumption for simplicity.

<sup>2</sup>Our approach can also be applied to some other loss functions. For example, in logistic regression for  $y \in \{\pm 1\}$ , when the loss is  $-\frac{1}{m} \sum_{l=1}^m \log [\text{sigmoid}(y_l f(x_l))]$ , we can bound the error  $\mathbb{E} [\text{sigmoid}(y f(x)) - 1]^2$ .

<sup>3</sup>Note that even though ReLU is not differentiable, we can still use its sub-gradient by defining  $\sigma'(u) = \mathbb{I}[u > 0]$ .

### 3.1 First order condition

In this section, we will rewrite the set of first order conditions for minimizing the empirical loss  $L$ . This rewriting motivates the direction of our later analysis. More specifically, the gradient of the empirical loss w.r.t.  $w_k$  is

$$\frac{\partial L}{\partial w_k} = \frac{1}{m} \sum_{l=1}^m (f(x_l) - y_l) v_k \sigma'(w_k^\top x_l) x_l, \quad (6)$$

for all  $k = 1, \dots, n$ . We will express this collection of gradient equations using matrix notation. Define the “extended feature matrix” as

$$D = \begin{pmatrix} v_1 \sigma'(w_1^\top x_1) x_1 & \cdots & v_1 \sigma'(w_1^\top x_m) x_m \\ \vdots & \ddots & \vdots \\ v_k \sigma'(w_k^\top x_1) x_1 & \cdots & v_k \sigma'(w_k^\top x_m) x_m \\ \vdots & \ddots & \vdots \\ v_n \sigma'(w_n^\top x_1) x_1 & \cdots & v_n \sigma'(w_n^\top x_m) x_m \end{pmatrix},$$

and the residual as

$$r = \frac{1}{m} (f(x_1) - y_1, \dots, f(x_m) - y_m)^\top.$$

Then we have

$$\frac{\partial L}{\partial W} := \left( \frac{\partial L}{\partial w_1}^\top, \dots, \frac{\partial L}{\partial w_n}^\top \right)^\top = D r. \quad (7)$$

A stationary point has zero gradient, so if  $D \in \mathbb{R}^{dn \times m}$  has full column rank, then immediately  $r = 0$ , *i.e.*, it is actually a global optimal. Since  $nd > m$  is necessary for  $D$  to have full column rank, we assume this throughout the paper.

However, in practice, we will not have the gradient being exactly zero because, *e.g.*, we stop the algorithm in finite steps or because we use stochastic gradient descent (SGD). In other words, typically we only have  $\|\partial L / \partial W\| \leq \epsilon$ , and  $D$  being full rank is insufficient since small gradient can still lead to large loss. More specifically, let  $s_m(D)$  be the minimum singular value of  $D$ , we have

$$\|r\| \leq \frac{1}{s_m(D)} \left\| \frac{\partial L}{\partial W} \right\|. \quad (8)$$

We can see that  $s_m(D)$  needs to be large enough for the residual to be small. Thus it is important to identify conditions to lower bound  $s_m(D)$  away from zero, which will be the focus of the paper.

### 3.2 Spectrum decay of activation kernel

We will later show that  $s_m(D)$  is related to the decay rate of the kernel spectrum associated with the activation function. More specifically, for an activation function  $\sigma(w^\top x)$ , we can define the following kernel function

$$g(x, y) = \mathbb{E}_w [\sigma'(w^\top x) \sigma'(w^\top y) \langle x, y \rangle] \quad (9)$$

where  $\mathbb{E}_w$  is over  $w$  uniformly distributed on a sphere.

In particular, for ReLU, the kernel has a closed form

$$g(x, y) = \left( \frac{1}{2} - \frac{\arccos \langle x, y \rangle}{2\pi} \right) \langle x, y \rangle. \quad (10)$$

In fact, it is a dot-product kernel and its spectrum can be obtained through spherical harmonic decomposition:

$$g(x, y) = \sum_{u=1}^{\infty} \gamma_u \phi_u(x) \phi_u(y), \quad (11)$$

where the eigenvalues are ordered  $\gamma_1 \geq \dots \geq \gamma_m \geq \dots \geq 0$  and the bases  $\phi_u(x)$  are spherical harmonics. The  $m$ -th eigenvalue  $\gamma_m$  will be related to  $s_m(D)$ .

For each spherical harmonic of order  $t$ , there are  $N(d, t) = \frac{2t+d-2}{t} \binom{t+d-3}{t-1}$  basis functions sharing the same eigenvalue. Therefore, the spectrum has a step like shape where each step is of length  $N(d, t)$ . Especially, for high dimensional input  $x$ , the number of such basis functions with large eigenvalues can be very large. Figure 1




Figure 1: The spectrum decay of the kernel associated with ReLU. We set  $d = 1500$ . It decays slower than  $O(1/m)$  for a large range of  $m$ .




Figure 2: The spectrum of a Gram matrix concentrates around the spherical harmonic spectrum of the kernel.

illustrates the spectrum of the kernel for  $d = 1500$ , and it is about  $\Omega(m^{-1})$  for a large range of  $m$ . For more details about the decomposition, please refer to Appendix A.

Such step like shape also appears in the Gram matrix associated with the kernel. Figure 2 compares the spectra of the kernel of  $d = 15$  and the corresponding Gram matrix with  $m = 3000$ . We can see the spectrum of the Gram matrix closely resembles that of the kernel. Such concentration phenomenon underlies the reason why the spectrum of  $D^\top D$  is closely related to the corresponding kernel.

### 3.3 Weight discrepancy

Another key factor in the analysis is the diversity of the unit weights, measured by the classic notion of geometric discrepancy [Bilyk and Lacey, 2015]. Given a set of  $n$  points  $W = \{w_k\}_{k=1}^n$  on the unit sphere  $\mathbb{S}^{d-1}$ , the discrepancy of  $W$  w.r.t. a measurable set  $S \subseteq \mathbb{S}^{d-1}$  is defined as

$$\text{dsp}(W, S) = \frac{1}{n} |W \cap S| - A(S), \quad (12)$$

where  $A(S)$  is the normalized area of  $S$  (i.e., the area of the whole sphere  $A(\mathbb{S}^{d-1}) = 1$ ).  $\text{dsp}(W, S)$  quantifies the difference between the empirical measure of  $S$  induced by  $W$  and the measure of  $S$  induced by a uniform distribution.

By defining a collection  $\mathcal{S}$  of such sets, we can summarize the difference in the empirical measure induced by  $W$  versus the uniform distribution over the sphere. More specifically, we will focus on the set of slices, each defined by a pair of inputs  $x, y \in \mathbb{S}^{d-1}$ , i.e.,

$$\begin{aligned} \mathcal{S} &= \{S_{xy} : x, y \in \mathbb{S}^{d-1}\}, \text{ where} \\ S_{xy} &= \{w \in \mathbb{S}^{d-1} : w^\top x \geq 0, w^\top y \geq 0\}. \end{aligned} \quad (13)$$

Essentially, each  $S_{xy}$  defines a slice-shaped area on the sphere which is carved out by the two half spaces  $w^\top x \geq 0$  and  $w^\top y \geq 0$ .

Based on the collection  $\mathcal{S}$ , we can define two discrepancy measures relevant to ReLU units.  $L_\infty$  discrepancy of  $W$  w.r.t.  $\mathcal{S}$  is defined as

$$L_\infty(W, \mathcal{S}) = \sup_{S \in \mathcal{S}} |\text{dsp}(W, S)|, \quad (14)$$

and the  $L_2$  discrepancy as

$$L_2(W, \mathcal{S}) = \sqrt{\mathbb{E}_{x,y} \text{dsp}(W, S_{xy})^2} \quad (15)$$

where the expectation is taken over  $x, y$  uniformly on the sphere. We use  $L_\infty(W)$  and  $L_2(W)$  as their shorthands. Both discrepancies measure how diverse the points  $W$  are. The more diverse the points, the smaller the discrepancy.

For our analysis, we slightly generalize the discrepancy for  $w_k$ 's not on unit sphere, by setting

$$\text{dsp}(\{w_k\}_k, S) = \text{dsp}(\{w_k/\|w_k\|\}_k, S). \quad (16)$$

## 4 Main results

Our main result is a bound on the training loss and the generalization error, assuming sufficiently large  $n, d$ . To state the theorem, first recall that  $\beta \in (0, 1)$  is the decay exponent of the spectrum of the activation kernel in (10), that is,  $\gamma_m$  is the  $m$ -th eigenvalue of the kernel and satisfies  $\gamma_m = \Omega(m^{-\beta})$ . Also recall that  $\mathcal{F}_W$  denote the set of feasible values of  $W$ .

**Theorem 1** (Main, simplified). *Let  $\delta, \delta' \in (0, 1)$ . If*

$$n = \tilde{\Omega}(m^\beta), \quad d = \tilde{\Omega}(m^\beta),$$

*then there exists a set  $\mathcal{G}_W \subseteq \mathcal{F}_W$  which takes up  $1 - \delta'$  fraction of measure of  $\mathcal{F}_W$ , such that with probability at least  $1 - cm^{-\log m} - \delta$  the following holds. For any  $W \in \mathcal{G}_W$  and any  $v \in \{-1, 1\}^n$ , we have*

$$\begin{aligned} \frac{1}{2m} \sum_{\ell=1}^m (f(x_\ell; v, W) - y_\ell)^2 &\leq c \left\| \frac{\partial L}{\partial W} \right\|^2, \\ \frac{1}{2} \mathbb{E}(f(x; v, W) - y)^2 &\leq c \left\| \frac{\partial L}{\partial W} \right\|^2 + c'(C_W + Y)^2 \sqrt{\frac{1}{m} \log \frac{1}{\delta}}. \end{aligned}$$

Here  $\tilde{\Omega}$  hides logarithmic terms  $\log m \log \frac{1}{\delta} \log \frac{1}{\delta'}$ .

**Remark 1.** Intuitively, the theorem means that when  $n, d$  are sufficiently large, for most feasible weights  $W$ , the training loss is in the order of the square norm of the gradient, and the generalization error has an additional term  $\tilde{O}(1/\sqrt{m})$ . In particular, when we obtain a solution  $W \in \mathcal{G}_W$  with gradient  $\|\partial L/\partial W\|^2 \leq \epsilon$ , the training loss is bounded by  $O(\epsilon)$ , and the generalization error is  $\tilde{O}(\epsilon + 1/\sqrt{m})$ . So neural networks trained with sufficiently many data points to a solution with small gradient in  $\mathcal{G}_W$  generalize well. This essentially means that although non-convex, the loss function over this set is well behaved, and thus learning is not difficult.

Note that an immediate corollary is that any critical point in  $\mathcal{G}_W$  is global optimum.

Furthermore, a randomly sampled set of weights  $W$  are likely to fall into this set. This suggests a reason for the practical success of training with random initialization: after initialization, the parameters w.h.p. fall into the set, then would stay inside during training, and finally get to a point with small gradient, which by our analysis, has small error.

**Remark 2.** An important feature about our result is that the set  $\mathcal{G}_W$  has a simple explicit form:

$$\mathcal{G}_W = \left\{ W \in \mathcal{F}_W : (L_2(W))^2 \leq c_g \left( \sqrt{\frac{\log d}{nd} \log \frac{1}{\delta'}} + \frac{1}{n} \log \frac{1}{\delta'} \right) \right\} \quad (17)$$

where  $c_g > 0$  is a universal constant. Furthermore,  $(L_2(W))^2$  has a simple closed form (See Theorem 7). Therefore, it is possible to directly check if a solution  $W$  is in  $\mathcal{G}_W$ , or design regularization that make sure  $W$  stays in the set  $\mathcal{G}_W$ .

**Remark 3.** The above theorem is a special case of the following more general result.

**Theorem 2** (Main, general). *Let  $\delta \in (0, 1)$ . With probability  $\geq 1 - m \exp(-m\gamma_m/8) - 2m^2 \exp(-4 \log^2 d) - \delta$ , the following holds. For any  $\xi, \eta > 0$  and any  $W$  with  $L_2(W) = \tilde{O}(n^{-\xi} d^{-\eta})$  such that*

$$m^\beta \leq \tilde{O} \left\{ d^{(1+\eta)/2} n^{\xi/2} m^{1/4}, d^{1/2} m^{1/2}, n^\xi d^{1/2+\eta} \right\}, \quad (18)$$

and any  $v \in \{-1, 1\}^n$ , we have

$$\begin{aligned} \frac{1}{2m} \sum_{\ell=1}^m (f(x_\ell; v, W) - y_\ell)^2 &\leq \frac{cm^\beta}{n} \left\| \frac{\partial L}{\partial W} \right\|^2, \\ \frac{1}{2} \mathbb{E}(f(x; v, W) - y)^2 &\leq \frac{cm^\beta}{n} \left\| \frac{\partial L}{\partial W} \right\|^2 + c'(C_W + Y)^2 \sqrt{\frac{1}{m} \log \frac{1}{\delta}}. \end{aligned}$$

Here  $\tilde{\Omega}$  hides logarithmic terms  $\log m \log \frac{1}{\delta} \log \frac{1}{\delta}$ .

The theorem shows that when the weights are diverse (i.e., with good discrepancy  $L_2(W)$  so that  $\xi$  and  $\eta$  are sufficiently large), the training loss is proportional to  $m^\beta/n$  times the squared norm of the gradient, and the generalization error has an additional term  $\tilde{O}(1/\sqrt{m})$ .

To obtain Theorem 1 from this more general result, we first note that Lemma 8 proves most  $W$  falls into  $\mathcal{G}_W$ , so it is sufficient to choose  $n, d$  large enough to guarantee the condition (18). Setting  $n = \tilde{O}(m^\beta)$  and  $d = \tilde{O}(m^\beta)$  satisfies the condition with  $\xi = \eta = 1/4$ , and thus leads to Theorem 1. Clearly, there exist some other options. For example,  $d = \tilde{\Omega}(m^{2\beta-1})$  and  $n = \tilde{\Omega}(m^{3-2\beta})$ , which matches the empirical observation that overspecified networks with large  $n$  are easier for the optimization.

## 5 Analysis roadmap

Our key technical result is a lower bound on the smallest singular value of  $D$  based on the spectrum of the activation kernel defined in (10) and the discrepancy of the weights defined in (15). Once the lower bound is obtained, we can use (8) to bound the training loss, and use Rademacher complexity to bound the generalization error.

**Theorem 3** (Smallest singular value). *With probability  $\geq 1 - m \exp(-m\gamma_m/8) - 2m^2 \exp(-4 \log^2 d) - \delta$ , the following holds. For any  $\xi, \eta > 0$ , any  $W$  with  $L_2(W) = \tilde{O}(n^{-\xi} d^{-\eta})$ , and any  $v \in \{-1, 1\}^n$ ,*

$$s_m(D)^2 \geq \Omega(nm^{1-\beta}) - \tilde{O}\left(\frac{n^{1-\xi/2} m^{3/4}}{d^{(1+\eta)/2}}\right) - \tilde{O}\left(\frac{nm^{1/2}}{d^{1/2}}\right) - \tilde{O}\left(\frac{n^{1-\xi} m}{d^{1/2+\eta}}\right).$$

Here the notation  $\tilde{\Omega}$  hides logarithmic terms  $\log d \log \frac{1}{\delta}$ .

The theorem is stated in its general form. It bounds the smallest singular value in terms of the  $n, d, m$  and two parameters  $\xi, \eta$  quantifying how large  $L_2(W)$  is. It is instructive to consider the special case when  $n = \tilde{\Omega}(m^\beta)$ ,  $d = \tilde{\Omega}(m^\beta)$ , and  $\xi = \eta = 1/4$ , which corresponds to Theorem 1). In this case, with probability at least  $1 - cm^{-\log m} - \delta$ ,  $s_m(D)^2 \geq cm$  for some constant  $c > 0$ . It is clear that the singular value is large and bounded away from zero.

It is interesting to compare the theorem to the results in [Soudry and Carmon, 2016], which shows that  $D$  is full rank with probability one under small perturbations. However, full-rankness alone is not sufficient since its smallest singular value could be extremely small leading to possibly huge training loss. Instead, we directly bound the smallest singular value and relate it to the activation and the diversity of the weights.

**Intuition.** Here we describe the high level intuition for bounding the minimum singular value. It is necessarily connected to the activation function and the diversity of the weights. For example, if  $\sigma'(t)$  is very small for all  $t$ , then the smallest singular value is expected to be very small. For the weights, if  $d < m$  (the interesting case) and all  $w_k$ 's are the same, then  $D$  cannot have rank  $m$ . If  $w_k$ 's are very similar to each other, then one would expect the smallest singular value to be very small or even zero. Therefore, some notion of diversity of the weights is needed.

The analysis begins by considering the matrix  $G_n = D^\top D/n$ . It suffices to bound  $\lambda_m(G_n)$ , the  $m$ -th (and the smallest) eigenvalue of  $G_n$ . To do so, we introduce a matrix  $G$  whose entries  $G(i, j) = \mathbb{E}_w[G_n(i, j)]$  where the expectation  $\mathbb{E}_w$  is taken assuming  $w_k$ 's are uniformly random on the unit sphere. The intuition is that when  $w$  is uniformly distributed,  $\sigma'(w^\top x)$  is most independent of the actual value of the  $x$ , and the matrix  $D$  should have the highest chance of having large smallest singular value. We will introduce  $G$  as an intermediate quantity and

subsequently bound the spectral difference between  $G_n$  and  $G$ . Roughly speaking, we break the proof into two steps

$$\lambda_m(G_n) \geq \underbrace{\lambda_m(G)}_{\text{I. ideal spectrum}} - \underbrace{\|G - G_n\|}_{\text{II. discrepancy}}$$

where  $\|G - G_n\|$  is the spectral norm of the difference.

For the first term in the lower bound, we observe that  $G$  has a particular nice form:  $G(i, j) = g(x_i, x_j)$ , the kernel defined in (10). This allows us to apply the eigendecomposition of the kernel and positive definite matrix concentration inequality to bound  $\lambda_m(G)$ , which turns out to be around  $m\gamma_m/2$ .

For the second term, when  $w_k$ 's are indeed from the uniform distribution over the sphere, this can be bounded by concentration bounds. It turns out that when  $w_k$ 's are not too far away from that, it is still possible to do so. Therefore, we use the geometric discrepancy to measure the diversity of the weights, and show that when they are sufficiently diverse,  $\|G - G_n\|$  is small. In particular, the entries in  $G - G_n$  can be viewed as the kernel of some U-statistics, hence concentration bounds can be applied. The expected U-statistics turns out to be the  $(L_2(W))^2$ , which has a closed form and can be shown to be small.

**Outline.** Theorem 3 is proved in Section 6,  $L_2(W)$  and  $\mathcal{G}_W$  are characterized in Section 6.2, and the proof sketch of Theorem 1 and Theorem 2 is provided in Section 7. We describe the proof sketch for the lemmas and provide the remaining proofs in the appendix.

## 6 Bounding the smallest singular value

Theorem 3 can be obtained from the following technical lemma.

**Lemma 4.** *With probability  $\geq 1 - m \exp(-m\gamma_m/8) - 2m^2 \exp(-4 \log^2 d) - \delta$ , we have*

$$s_m(D)^2 \geq nm\gamma_m/2 - cn\rho(W), \quad (19)$$

where

$$\rho(W) = \frac{\log d}{\sqrt{d}} \sqrt{L_\infty(W)L_2(W)} m \left( \frac{4}{m} \log \frac{1}{\delta} \right)^{1/4} + \frac{\log d}{\sqrt{d}} mL_\infty(W) \sqrt{\frac{4}{3m} \log \frac{1}{\delta}} + \frac{\log d}{\sqrt{d}} mL_2(W) + L_\infty(W). \quad (20)$$

*Proof of Theorem 3.* First,  $|\text{dsp}(W, S)| \leq 2$  for any set  $W$  and slice  $S$ , so by definition  $|L_\infty(W)| \leq 2$ . Next, By the assumption in the theorem,  $L_2(W) = \tilde{O}(n^{-\xi} d^{-\eta})$ . Plugging these into Lemma 4 completes the proof.  $\square$

Lemma 4 is meaningful only when  $cn\rho(W)$  is small compared to  $nm\gamma_m/2$ . This requires  $L_2(W)$  to be sufficiently small. In the following we will first provide the proof sketch of Lemma 4, and then bound that  $L_2(W)$  in Section 6.2.

### 6.1 Proof of Lemma 4

To prove Lemma 4, it is sufficient to bound the smallest eigenvalue of  $G_n = D^\top D/n$ . Note that  $v_k \in \{-1, 1\}$ , so  $v_k^2 = 1$ , and thus the  $(i, j)$ -th entry of  $G_n$  is

$$G_n(i, j) = \frac{1}{n} \sum_{k=1}^n \sigma'(w_k^\top x_i) \sigma'(w_k^\top x_j) \langle x_i, x_j \rangle. \quad (21)$$

For ReLU,  $\sigma'(w^\top x)$  does not depend on the norm of  $w$  so without loss of generality, we assume  $\|w\| = 1$ . Consider a related matrix  $G$  whose  $(i, j)$ -th entry is defined as

$$G(i, j) = \mathbb{E}_w [\sigma'(w^\top x_i) \sigma'(w^\top x_j) \langle x_i, x_j \rangle]. \quad (22)$$

where  $w$  is a random variable distributed uniformly over the unit sphere. Since  $\sigma'(w^\top x) = \mathbb{I}[w^\top x \geq 0]$ , we have a closed form expression for  $G(i, j)$ :

$$\begin{aligned} G(i, j) &= \mathbb{E}_w [\mathbb{I}(w^\top x_i \geq 0) \mathbb{I}(w^\top x_j \geq 0)] \langle x_i, x_j \rangle \\ &= \left( \frac{1}{2} - \frac{\arccos \langle x_i, x_j \rangle}{2\pi} \right) \langle x_i, x_j \rangle. \end{aligned} \quad (23)$$



Note that  $G(i, j) = g(x_i, x_j)$  where  $g$  is the kernel defined in (10). This allows us to reason about the eigenspectrum of  $G$ , denoted as  $\lambda_1(G) \geq \dots \geq \lambda_m(G)$ .

Therefore, our strategy is to first bound  $\lambda_m(G)$  in Lemma 5 and then bound  $|\lambda_m(G) - \lambda_m(G_n)|$  in Lemma 6. Combining the two immediately leads to Lemma 4.

**First, consider  $\lambda_m(G)$ .** We consider a truncated version of spherical harmonic decomposition:

$$g^{[m]}(x_i, x_j) = \sum_{u=1}^m \gamma_u \phi_u(x_i) \phi_u(x_j)$$

and the corresponding matrix  $G^{[m]}$ . On one hand, it is clear that  $\lambda_m(G) \geq \lambda_m(G^{[m]})$ . On the other hand,  $G^{[m]} = AA^\top$  where  $A$  is a random matrix whose rows are

$$A^i := [\sqrt{\gamma_1} \phi_1(x_i), \dots, \sqrt{\gamma_m} \phi_m(x_i)].$$

Next, we bound  $\lambda_m(G^{[m]})$  by matrix Chernoff bound [Tropp, 2012], and it is better than previous work [Braun, 2006]. This leads to the following lemma.

**Lemma 5.** *With probability at least  $1 - m \exp(-m\gamma_m/8)$ ,*

$$\lambda_m(G) \geq m\gamma_m/2.$$

**Next, bound  $|\lambda_m(G) - \lambda_m(G_n)|$ .** By Weyl's theorem, this is bounded by  $\|G - G_n\|$ . To simplify the notation, denote

$$E_{i,j} = \mathbb{E}_w[\sigma'(w^\top x_i) \sigma'(w^\top x_j)] - \frac{1}{n} \sum_{k=1}^n \sigma'(w_k^\top x_i) \sigma'(w_k^\top x_j).$$

Then  $G(i, j) - G_n(i, j) = \langle x_i, x_j \rangle E_{i,j}$ , and thus

$$\begin{aligned} & \|G - G_n\| \\ &= \sup_{\|\alpha\|=1} \sum_{i,j} \alpha_i \alpha_j \langle x_i, x_j \rangle E_{i,j} \\ &\leq \sup_{\|\alpha\|=1} \sqrt{\sum_{i \neq j} \alpha_i^2 \alpha_j^2} \sqrt{\sum_{i \neq j} \langle x_i, x_j \rangle^2 E_{i,j}^2} + \max_i |E_{ii}| \\ &\leq \frac{c \log d}{\sqrt{d}} \sqrt{\sum_{i \neq j} E_{i,j}^2} + \max_i |E_{ii}| \end{aligned} \tag{24}$$

where the last inequality holds with high probability since  $x_i$ 's are uniform over the unit sphere and thus we can apply sub-gaussian concentration bounds.

Note that  $\sum_{i \neq j} E_{i,j}^2 / (m(m-1))$  is a U-statistics where the summands are dependent and typical concentration inequality for *i.i.d.* entries does not apply. Instead we use a Bernstein inequality for U-statistics [Peel et al., 2010] to show that with probability at least  $1 - \delta$ , it is bounded by

$$\mathbb{E}_{\{x_1, x_2\}} E_{12}^2 + \sqrt{\frac{4\Sigma^2}{m} \log \frac{1}{\delta}} + \frac{4B^2}{3m} \log \frac{1}{\delta} \tag{25}$$

where  $B = \max_i |E_{ii}|$  and  $\Sigma^2 = \mathbb{E}[E_{12}^4] - (\mathbb{E}[E_{12}^2])^2$ .

The key observation is that the quantities in the above lemma are related to discrepancy:

$$\max_{i,j} E_{i,j} \leq L_\infty(W), \tag{26}$$

$$\mathbb{E}_{x_1, x_2} [E_{12}^2] = (L_2(W))^2, \tag{27}$$

$$\Sigma^2 \leq (L_2(W) L_\infty(W))^2. \tag{28}$$

This is because  $\sigma'(w^\top x_i)\sigma'(w^\top x_j) = \mathbb{I}[w \in S_{x_i x_j}]$  and thus

$$\begin{aligned} E_{i,j} &= \mathbb{E}_w \mathbb{I}[w \in S_{x_i x_j}] - \frac{1}{n} \sum_{k=1}^n \mathbb{I}[w_k \in S_{x_i x_j}] \\ &= A(S_{x_i x_j}) - \frac{1}{n} |W \cap S_{x_i x_j}| \\ &= -\text{dsp}(W, S_{x_i x_j}). \end{aligned}$$

Plugging (26)-(28) into (25) and (24), we have

**Lemma 6.** *The following inequality holds with probability at least  $1 - 2m^2 \exp(-\log^2 d) - \delta$ ,*

$$\|G_n - G\| \leq c\rho(W) \quad (29)$$

where  $\rho(W)$  is as defined in Lemma 4.

## 6.2 Characterizing the discrepancy

In this subsection, we present a bound for  $L_2(W)$  and show that the  $\mathcal{G}_W$  defined in the following covers most  $W$ 's. Recall that

$$\mathcal{G}_W = \left\{ W : (L_2(W))^2 \leq c_g \left( \sqrt{\frac{\log d}{nd} \log \frac{1}{\delta'}} + \frac{1}{n} \log \frac{1}{\delta'} \right) \right\}$$

for  $0 < \delta' < 1$  and a proper constant  $c_g > 0$ . The constant  $c_g$  is the constant in Lemma 8.  $\delta'$  will be clear from the context where  $\mathcal{G}_W$  is used.

First we provide a closed form for  $L_2$  discrepancy of slices defined in (13). The proof is provided in the appendix.

**Theorem 7.** *Suppose  $W = \{w_i\}_{i=1}^n \subseteq \mathbb{S}^{d-1}$ .*

$$(L_2(W))^2 = \frac{1}{n^2} \sum_{i,j=1}^n k(w_i, w_j)^2 - \mathbb{E}_{u,v} [k(u, v)^2]$$

where  $\mathbb{E}_{u,v}$  is over  $u$  and  $v$  uniformly distributed on  $\mathbb{S}^{d-1}$  and the kernel  $k(\cdot, \cdot)$  is

$$k(u, v) = \frac{\pi - \arccos \langle u, v \rangle}{2\pi}.$$

The closed form is simple and intuitive. The kernel  $k(w_i, w_j)$  measures how similar two units are. The discrepancy is the difference between the average pairwise similarity and the expected one over uniform distribution.

Given the theorem, we now show that  $(L_2(W))^2$  can be small. We use the probabilistic method, *i.e.*, show that if  $w_k$ 's are sampled from  $\mathbb{S}^{d-1}$  uniformly at random, then with high probability  $W$  falls into  $\mathcal{G}_W$ . The key observation is that with random  $W$ ,  $(L_2(W))^2$  is the difference between a U-statistics and its expectation, which can be bounded by concentration inequalities. Formally,

**Lemma 8.** *There exists a constant  $c_g$ , such that for any  $0 < \delta' < 1$ , with probability at least  $1 - \delta'$  over  $W = \{w_i\}_{i=1}^n$  that are sampled from the unit sphere uniformly at random,*

$$(L_2(W))^2 \leq c_g \left( \sqrt{\frac{\log d}{nd} \log \frac{1}{\delta'}} + \frac{1}{n} \log \frac{1}{\delta'} \right).$$

Alternatively, the theorem means that  $\mathcal{G}_W$  defined in (17) covers most  $W$ 's. This is because  $L_2(W)$  is independent of the length of  $w_k$ 's, it is sufficient to show that  $L_2(W)$  is small for  $W \in \mathcal{G}_W \cap \mathbb{S}^{d-1}$ .

## 7 Final bound on generalization error

Here we provide the proof sketch of Theorem 2 and Theorem 1. More details of the proof are in Appendix F.

First, we prove Theorem 2. Suppose a solution  $W$  satisfies the assumption and has small gradient  $\|\partial L/\partial W\|$ . Using (8), we have  $\|r\| \leq \|\partial L/\partial W\|/s_m(D)$ . By Theorem 3 and the assumption in Theorem 2, with high probability  $s_m^2(D) = \Omega(nm^{1-\beta})$ . This implies the training loss is

$$\frac{1}{m} \sum_l (f(x_l) - y_l)^2 = m \|r\|^2 \leq \frac{cm^\beta}{n} \left\| \frac{\partial L}{\partial W} \right\|^2.$$

The generalization error can then be derived using McDiamid's inequality and Rademacher complexity. First, we need an upper bound on the difference of the loss for two data points for the McDiamid's inequality. Since  $\|x\|_2 \leq 1$  and  $\sum_k \|w_k\|_2 \leq C_W$ , we have  $|f| \leq C_W$ . Thus

$$\begin{aligned} & |\ell(y, f(x)) - \ell(y', f(x'))| \\ & \leq \frac{1}{2} \max \{ (y - f(x))^2, (y' - f(x'))^2 \} \leq Y^2 + C_W^2 \end{aligned}$$

where in the last inequality we use the fact that the true function  $|y| \leq Y$ . Next, we use the composition rules to compute the Rademacher complexity. Since the complexity of linear functions  $\{w^\top x : \|w\|_2 \leq b_W, \|x\|_2 \leq 1\}$  is  $b_W/\sqrt{m}$  and  $\sigma(\cdot)$  is 1-Lipschitz, and  $\sum_k \|w_k\|_2 \leq C_W$ , the complexity  $\mathcal{R}_m(\mathcal{F}) \leq C_W/\sqrt{m}$ . Composing it with the loss function, and applying the bound in [Bartlett and Mendelson, 2002], we get the final generalization bound.

To obtain Theorem 1 from the more general Theorem 2, we first note that Lemma 8 proves most  $W$  falls into  $\mathcal{G}_W$ , and setting  $n = \tilde{O}(m^\beta)$  and  $d = \tilde{O}(m^\beta)$  makes sure that any  $W \in \mathcal{G}_W$  has  $L_2(W) = \tilde{O}(n^{-1/4}d^{-1/4})$  satisfying the condition (18).

## 8 Discussions

In this section, we discuss and remark on further considerations and possible extensions of our current analysis.

### 8.1 Other loss functions

Currently, our analysis is tied to the least squares loss  $\ell(y, f(x)) = \frac{1}{2} (y - f(x))^2$ . It is fairly straightforward to generalize it to any strongly convex loss function, such as logistic loss. Note that the final objective function is *not* convex due to the non-convexity in neural networks, but most loss functions used in practice are strongly convex w.r.t.  $f(x)$ . Under the new setting, the residual is then

$$r = \frac{1}{m} (\ell'(y_1, f(x_1)), \dots, \ell'(y_m, f(x_m)))^\top.$$

According to our analysis, the norm of the residual  $\|r\|$  will be bounded. This in turn implies each individual  $\ell'(y_l, f(x_l))$  will be small. Since the loss function  $\ell(y, f(x))$  is strongly convex, the loss itself will be small.

### 8.2 Other activation functions

We can consider a family of activation functions of the form  $\sigma(u) = \max\{u, 0\}^t$ , *i.e.*, rectified polynomials [Cho and Saul, 2009, Krotov and Hopfield, 2016]. This requires two modifications to the analysis.

One is the corresponding kernel  $k(x, y) = \mathbb{E}_w [\sigma'(w^\top x) \sigma'(w^\top y)]$  and  $g(x, y) = k(x, y) \langle x, y \rangle$ . When the input distribution is uniform, we can also compute the kernels in closed form as shown in [Cho and Saul, 2009]:

$$k_t(x, y) = \frac{J_{t-1}(\theta)}{2\pi} \tag{30}$$

where

$$J_t(\theta) = (-1)^t (\sin \theta)^{2t+1} \left( \frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \right)^t \left( \frac{\pi - \theta}{\sin \theta} \right), \tag{31}$$

and  $\theta = \arccos \langle x, y \rangle$ . Note that the subscript is  $t - 1$  in (30) because we are computing on the derivative  $\sigma'(u)$ .

Examples for the first few  $t$  are listed as follows.

$$J_0(\theta) = \pi - \theta \quad (32)$$

$$J_1(\theta) = \sin \theta + (\pi - \theta) \cos \theta \quad (33)$$

$$J_2(\theta) = 3 \sin \theta \cos \theta + (\pi - \theta)(1 + 2 \cos^2 \theta) \quad (34)$$

Larger  $t$  corresponds to more nonlinear activation functions and leads to slower decaying spectrum since there are more high frequency components.

We also need to change the definition of the discrepancy to accommodate the new kernels. Let

$$\begin{aligned} (L_2(W))^2 &= \mathbb{E}_{x_i, x_j} \left[ \mathbb{E}_w [\sigma'(w^\top x_i) \sigma'(w^\top x_j)] - \frac{1}{n} \sum_{k=1}^n \sigma'(w_k^\top x_i) \sigma'(w_k^\top x_j) \right]^2 \\ &= \frac{1}{n^2} \sum_{i,j=1}^n k(w_i, w_j)^2 - \mathbb{E}_{u,v} [k(u, v)^2]. \end{aligned} \quad (35)$$

Therefore, the discrepancy is affected by how the kernels change due to change in activation functions.

The other modification is on the Rademacher complexity. Since the derivative  $\sigma'(u) = t \max\{u, 0\}^{t-1}$ , there is an additional factor of  $t$  in front of the complexity. That is, larger  $t$  leads to higher Rademacher complexity.

In summary, the best parameter  $t$  depends on the balance between the two conflicting effects. On one hand, larger  $t$  corresponds to slower decaying spectrum and makes the minimum singular value more likely to be larger. On the other hand, smaller  $t$  leads to better generalization since the Rademacher complexity is smaller.

### 8.3 (Sub)gradient of the activation function

Throughout this paper, we have used one particular subgradient for the ReLU activation function:  $\mathbb{I}[u > 0]$ . At the point  $u = 0$ , there are many other valid subgradients as long as its value is between 0 and 1. However, our analysis is not restricted to this particular subgradient. First of all, all the subgradients only differ at one point  $u = 0$ , which is of probability zero. Second, our analysis is probabilistic in nature. The first term in Lemma 4 is the expectation over  $W$ , which is insensitive to the probability zero event  $u = 0$ . The second term in Lemma 4 is related to  $L_2(W)$ , which is again expectation over all possible data, thus insensitive to the difference.

In summary, though for some  $W \in \mathcal{G}_W$  the loss is not differentiable, one can define  $\partial L / \partial W$  by using subgradients of ReLU  $\sigma$  as follows:

$$\sigma'(x) = \begin{cases} 0, & x < 0 \\ c, & x = 0 \\ 1, & x > 0 \end{cases} \quad (36)$$

for any  $c \in [0, 1]$ . Then under the conditions in our theorems, with high probability, for any  $W \in \mathcal{G}_W$  and any definition of  $\sigma'$  in (36), the guarantees hold.

Other activation functions such as rectified polynomials are differentiable and thus they do not have such issue.

### 8.4 Other input distribution

When the input distribution is not uniform, the spectrum of the kernel function defined in (9) will be different because the spherical harmonic bases are defined with respect to the input distribution. To ensure the spectrum decays slowly, we need to find a corresponding distribution of  $W$  that “matches” the input distribution.

We provide some intuitions in finding such “matching” distribution. Suppose the input distribution is uniform on the set  $E \in \mathbb{S}^{d-1}$ , if a hyperplane whose normal is  $w$  does not “cut through” the set, then for all data points, they have the same sign  $\mathbb{I}[w^\top x > 0]$ . This will likely lead to rank deficiency in the extended feature matrix.

Therefore, we prefer  $W$  to split the data points as much as possible. One such distribution of  $W$  is uniform on the set  $F_E = \{w \in \mathbb{S}^{d-1} : \text{there exists } u \in E, \langle u, w \rangle = 0\}$ . For example, if  $E$  is the intersection of the positive orthant and the unit sphere,  $E = \{u \in \mathbb{S}^{d-1} : u_i \geq 0, \text{ for all } i \in [d]\}$ , then the corresponding set  $F_E$  is the whole sphere excluding  $E$  and  $-E$ .

Table 1: Comparison of minimum eigenvalues with uniform and “matching” distributions. Note that the “matching” distribution corresponds to larger minimum eigenvalue for different dimensions. However, the difference becomes negligible when the dimension increases.

$d$	4	5	6	7
uniform	$3.96 \times 10^{-4}$	0.0015	0.0032	0.0072
matching	<b><math>5.43 \times 10^{-4}</math></b>	<b>0.0017</b>	0.0032	0.0072

We have verified the phenomenon empirically. We have generated 3000 input data points uniform on the positive orthant  $E$ . We then compute the  $3000 \times 3000$  Gram matrix, where the  $(i, j)$ -th entry is  $\mathbb{E}_w [\sigma'(w^\top x_i) \sigma'(w^\top x_j) \langle x_i, x_j \rangle]$ . The expectation is approximated by sampling 100,000 independent  $w$ ’s and then averaging. We compare two distributions of  $W$ : 1) uniform on the whole unit sphere; 2) uniform on  $F_E$ .

In Table 1, we compare the minimum eigenvalues with the two distributions. The uniform distribution on  $F_E$  always leads to larger or the same minimum eigenvalues. However, as dimension increases, the difference becomes negligible. Note that the difference between the uniform distribution on the whole sphere and uniform on  $F_E$  becomes exponentially small when the dimension  $d$  increases, because the proportion of  $E$  and  $-E$  shrinks exponentially. This suggests that in high dimensions, uniform on the whole unit sphere is a reasonable distribution for  $W$ .

For a general input distribution  $P(x)$ , we can decompose it into small sets  $dx$  and on every set, the distribution is uniform with measure  $P(x)dx$ . Then every small sets corresponds to a distribution of  $W$ . The final distribution of  $W$  is the superposition of all such distributions, weighted by  $P(x)dx$ .

## 9 Numerical evaluation

In this section, we further investigate numerically the effects of gradient descent on the discrepancy and the effects of regularizing the weights using discrepancy measure.

### 9.1 Discrepancy and gradient descent

One limitation of the analysis is that we have not analyzed how to obtain a solution  $W \in \mathcal{G}_W$  with small gradient. The theoretical analysis of gradient descent is left for future work. Meanwhile we provide some numerical results supporting our claims.

Although the set  $\mathcal{G}_W$  contains most  $W$ ’s, it is still unclear whether the solutions given by gradient descent lie in the set. We design experiments to investigate this issue. The ground truth input data are of dimension  $d = 50$  and true function consists of  $n = 50$  units. We use networks of different  $n$  to learn the true function and perform SGD with batch size 100 and learning rate 0.1 for 5000 iterations. Figure 3 shows how  $(L_2(W))^2$  changes as a function of  $n$ . It is slightly worse than  $O(n^{-1})$  but scales better than  $O(n^{-1/2})$ , suggesting (stochastic) gradient descent outputs solutions with reasonable discrepancy.

### 9.2 Regularization

To reinforce solutions with small discrepancy, we propose a novel regularization term to minimize  $L_2$  discrepancy:

$$R(W) = \frac{1}{n(n-1)} \sum_{i \neq j}^n k(w_i, w_j)^2. \quad (37)$$

It is essentially  $L_2$  discrepancy without the constants.

To verify the effectiveness of the regularization term, we explore the relationship between the regularization and the minimum singular value. We first generate 20 random  $W$ ’s, all with  $n = 100$  and  $d = 100$ , and compute their discrepancy and singular values using  $m = 3000$ . Then we optimize  $R(W)$  and compare the quantities after optimization. The result is presented in Figure 4. We can see smaller regularization value corresponds to larger singular value.

We also conducts experiments to compare training and test errors with and without regularization. The ground truth data are of  $d = 100$  and  $n = 100$ . We learn the true function by SGD with learning rate 0.1, momentum 0.9




Figure 3: Discrepancy of  $W$  obtained after gradient descent. We perform gradient descent and compute discrepancy for the returned solution. The red curve corresponds to such solutions with different  $n$ . It scales similarly to the bound for uniform  $W$  as in Lemma 8.




Figure 4: Effect of regularization. The blue dots represent random weights and the red dots linked with dashed black lines represent weights optimized by minimizing  $R(w)$ . Smaller regularization values correspond to larger minimal singular values.

Table 2: Comparison of performance with/without regularization (all numbers are of unit  $10^{-5}$ ). The true function is generated with  $d = 100$  and  $n = 100$ . To learn the function, we use networks with different  $n$ .

	$n = 100$		$n = 150$	
	train	test	train	test
no-reg	15.42(5.86)	14.80(5.36)	1.79(0.45)	1.86(0.50)
reg	<b>11.32(1.77)</b>	<b>10.63(1.58)</b>	<b>1.07(0.84)</b>	<b>1.13(0.99)</b>
	$n = 200$		$n = 300$	
	train	test	train	test
no-reg	<b>0.38(0.27)</b>	<b>0.44(0.35)</b>	0.39(0.39)	0.44(0.40)
reg	0.50(0.51)	0.58(0.59)	<b>0.10(0.05)</b>	<b>0.12(0.07)</b>

and a total of 300,000 iterations. The regularization coefficients are chosen from  $\{1, 0.1, 0.01, 0.001\}$  and the best results are reported. We use neural networks of size  $n \in \{100, 150, 200, 300\}$  and for each  $n$  we repeat five times with different random seeds. The result is summarized in Table 2. Regularization leads to lower training and test errors for most settings. Even in the case where the un-regularized one performs better, the errors are all small enough (within the same range as standard deviation), suggesting the noise begins to dominate.

## 10 Conclusion

We have analyzed one-hidden-layer neural networks and identified novel conditions when they achieve small training and test errors despite the non-convexity of the loss function. The key factors are the spectrum of the kernel associated with the activation function and the diversity of the units measured by discrepancy.

Although we focus on a least-square loss function and uniform input distribution, the analysis technique can be readily extended to other loss function and input distributions. At the moment, our analysis is still limited in the sense that it is independent of the actual algorithm. In the future work, we will explore the interplay between the discrepancy and gradient descent. In addition, we will further investigate the issue of designing an algorithm that guarantees good discrepancy thus small errors, possibly in a way similar to [Ge et al., 2016] in low-rank recovery problems.

## Acknowledge

We thank Santosh Vempala, Lorenzo Rosasco and Jason Lee for valuable discussions. The research is supported in part by NSF/NIH BIGDATA 1R01GM108341, ONRN00014-15-1-2340, NSF IIS-1639792, NSF IIS-1218749, NSF CAREER IIS-1350983, Intel and NVIDIA, and by NSF grants CCF-0832797, CCF-1117309, CCF-1302518, DMS-1317308, Simons Investigator Award, and Simons Collaboration Grant.

## References

- [Barron, 1993] Barron, A. R. (1993). Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Trans. Inform. Theory*, 39(3):930–945.
- [Bartlett and Mendelson, 2002] Bartlett, P. L. and Mendelson, S. (2002). Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482.
- [Bilyk and Lacey, 2015] Bilyk, D. and Lacey, M. T. (2015). One bit sensing, discrepancy, and stolarsky principle. *arXiv preprint arXiv:1511.08452*.
- [Blum and Rivest, 1993] Blum, A. L. and Rivest, R. L. (1993). Training a 3-node neural network is np-complete. In *Machine learning: From theory to applications*.
- [Braun, 2006] Braun, M. (2006). Accurate error bounds for the eigenvalues of the kernel matrix. *JMLR*, 7:2303–2328.
- [Cho and Saul, 2009] Cho, Y. and Saul, L. K. (2009). Kernel methods for deep learning. In Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C., and Culotta, A., editors, *Advances in Neural Information Processing Systems 22*, pages 342–350.
- [Choromanska et al., 2015] Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B., and LeCun, Y. (2015). The loss surfaces of multilayer networks. In *AISTATS*.
- [Dauphin et al., 2014] Dauphin, Y. N., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., and Bengio, Y. (2014). Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Advances in neural information processing systems*.
- [Ge et al., 2015] Ge, R., Huang, F., Jin, C., and Yuan, Y. (2015). Escaping from saddle points – online stochastic gradient for tensor decomposition. *arXiv preprint arXiv:1503.02101*.
- [Ge et al., 2016] Ge, R., Lee, J. D., and Ma, T. (2016). Matrix completion has no spurious local minimum. *arXiv preprint arXiv:1605.07272*.
- [Janzamin et al., 2015] Janzamin, M., Sedghi, H., and Anandkumar, A. (2015). Beating the perils of non-convexity: Guaranteed training of neural networks using tensor methods. *CoRR abs/1506.08473*.
- [Kawaguchi, 2016] Kawaguchi, K. (2016). Deep learning without poor local minima. In *Advances in Neural Information Processing Systems (NIPS)*.
- [Krotov and Hopfield, 2016] Krotov, D. and Hopfield, J. J. (2016). Dense associative memory for pattern recognition. *CoRR*, abs/1606.01164.
- [Lee et al., 2016] Lee, J., Simchowitz, M., Jordan, M., and Recht, B. (2016). Gradient descent only converges to minimizers. In *Proceedings of the Annual Conference on Learning Theory (COLT)*.
- [Müller, 2012] Müller, C. (2012). *Analysis of spherical symmetries in Euclidean spaces*, volume 129. Springer Science & Business Media.
- [Peel et al., 2010] Peel, T., Anthoine, S., and Ralaivola, L. (2010). Empirical bernstein inequalities for u-statistics. In *Advances in Neural Information Processing Systems*, pages 1903–1911.

- [Rahimi and Recht, 2009] Rahimi, A. and Recht, B. (2009). Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *Neural Information Processing Systems*.
- [Soudry and Carmon, 2016] Soudry, D. and Carmon, Y. (2016). No bad local minima: Data independent training error guarantees for multilayer neural networks. *arXiv preprint arXiv:1605.08361*.
- [Tropp, 2012] Tropp, J. A. (2012). User-friendly tools for random matrices: An introduction. Technical report, DTIC Document.
- [Vershynin, 2010] Vershynin, R. (2010). Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*.
- [Yang et al., 2014] Yang, Z., Moczulski, M., Denil, M., de Freitas, N., Smola, A. J., Song, L., and Wang, Z. (2014). Deep fried convnets. *CoRR*, abs/1412.7149.



## A Spherical harmonic decomposition and kernel spectrum

Any function defined on the unit sphere has a spherical harmonic decomposition

$$g(x, y) = \sum_u \gamma_u \phi_u(x) \phi_u(y), \quad (38)$$

where  $\phi_u(x) : \mathbb{S}^{d-1} \mapsto \mathbb{R}$  is a spherical harmonic basis. Note that  $u = (t, j)$  is a multi-index: the first denotes the order of the basis and the latter denotes the index of bases with the same order.

For each order  $t$ , there are  $N(d, t) = \frac{2t+d-2}{t} \binom{t+d-3}{t-1}$  bases with the same coefficient. As a result, the spectrum  $\gamma_u$  sorted by magnitude has the step like shape where each step is of length  $N(d, t)$ .

To compute the coefficients, we use the Legendre harmonics [Müller, 2012] with the following property

$$P_{t,d}(\langle x, y \rangle) = \frac{1}{N(d, t)} \sum_{j=1}^{N(d, t)} \phi_{t,j}(x) \phi_{t,j}(y). \quad (39)$$

The spherical harmonics also form an orthonormal basis on the unit sphere:

$$\mathbb{E} [\phi_{l,i}(x) \phi_{m,j}(x)] = \frac{1}{|\mathbb{S}^{d-1}|} \int_{\mathbb{S}^{d-1}} \phi_{l,i}(x) \phi_{m,j}(x) dx = \delta_{lm} \delta_{ij}, \quad (40)$$

where  $|\mathbb{S}^{d-1}|$  denotes the surface area of the unit sphere.

Combining these properties, we can calculate the spectrum using

$$\gamma_{(t,j)} = \frac{|\mathbb{S}^{d-2}|}{|\mathbb{S}^{d-1}|} \int_{-1}^1 g(\xi) P_{t,d}(\xi) (1 - \xi^2)^{(d-3)/2} d\xi, \quad \text{for all } j \in [N(d, t)]. \quad (41)$$

## B Bounding $\lambda_m(G)$ using matrix concentration bound: Proof of Lemma 5

Recall that

$$g(x, y) = \sum_{u=1}^{\infty} \gamma_u \phi_u(x) \phi_u(y). \quad (42)$$

For an integer  $r > 0$ , define the truncated version of  $g$  and the corresponding residue as

$$g^{[r]}(x, y) = \sum_{u=1}^r \gamma_u \phi_u(x) \phi_u(y) \quad (43)$$

$$e^r(x, y) = g(x, y) - g^{[r]}(x, y) \quad (44)$$

and define the matrices

$$\begin{aligned} [G^{[r]}]_{i,j} &= g^{[r]}(x_i, x_j) \\ E^r &= G - G^{[r]}. \end{aligned} \quad (45)$$

**Lemma 9.** Let  $c_g = \max_x g(x, x)$  then with probability at least  $1 - m \exp\left(-\frac{m\gamma_m}{8c_g}\right)$ ,

$$\lambda_m(G^{[m]}) \geq m\gamma_m/2.$$

*Proof.* Define a matrix  $A$  whose rows are

$$A^i := [\sqrt{\gamma_1} \phi_1(x_i), \dots, \sqrt{\gamma_m} \phi_m(x_i)]$$

for  $1 \leq i \leq m$ . Define matrices

$$X_i = (A^i)^\top A^i.$$

Denote  $Y = \sum_{i=1}^m X_i$ . Then  $\lambda_m(\mathbb{E}Y) = m\gamma_m$  using the fact that  $\mathbb{E}[\phi_i(x)\phi_j(x)] = \delta_{ij}$ . Furthermore,  $X_i \succeq 0$  and

$$\|X_i\| \leq \text{tr}(X_i) = \sum_{u=1}^m \gamma_u \phi_u^2(x_i) \leq g(x_i, x_i) = c_g.$$

Therefore, matrix Chernoff bound (e.g., [Tropp, 2012]) gives

$$\Pr[\lambda_m(Y) \leq (1 - \epsilon)\lambda_m(\mathbb{E}Y)] \leq m \exp\left(- (1 - \epsilon)^2 \lambda_m(\mathbb{E}Y)/(2c_g)\right).$$

Choose  $\epsilon = 1/2$  and use the facts that  $G^{[m]} = AA^\top$ ,  $Y = A^\top A$  and  $\lambda_m(G^{[m]}) = \lambda_m(Y)$ , we finish the proof.  $\square$

*Proof of Lemma 5.* By Weyl's theorem and the fact that  $E^m$  is PSD,

$$\lambda_m(G) \geq \lambda_m(G^{[m]}) + \lambda_m(E^m) \geq \lambda_m(G^{[m]}).$$

Lemma 5 then follows from Lemma 9.  $\square$

## C Bounding the difference between $\lambda_m(G)$ and $\lambda_m(G_n)$ : Proof of Lemma 6

Using Weyl's theorem we have that

$$|\lambda_m(G_n) - \lambda_m(G)| \leq \|G_n - G\|. \quad (46)$$

We are going to give an upper bound on  $\|G_n - G\|$ :

$$\|G_n - G\| = \sup_{\|\alpha\|=1} \sum_{i,j=1}^m \alpha_i \alpha_j (x_i^\top x_j) E_{ij}, \quad (47)$$

$$\text{where } E_{i,j} = \frac{1}{n} \sum_{k=1}^n \sigma'(w_k^\top x_i) \sigma'(w_k^\top x_j) - \mathbb{E}_w[\sigma'(w^\top x_i) \sigma'(w^\top x_j)], \quad (48)$$

and the first expectation is taken over  $w$  uniformly on the sphere  $\mathbb{S}^{d-1}$ .

Our bound heavily relies on the inner products  $|\langle x_i, x_j \rangle|$  for all  $i \neq j$  being small enough. In the next lemma, we provide such a result for uniformly distributed data.

**Lemma 10** (Tail bound for spherical distribution). *If  $a$  and  $b$  are independent vectors uniformly distributed over the unit sphere  $\mathbb{S}^{d-1}$ , then there exists a constant  $c > 0$ , such that for any  $u > 0$ ,*

$$\Pr\left[|\langle a, b \rangle| \geq \frac{cu}{\sqrt{d}}\right] \leq 2e^{-u^2}.$$

*Proof.* Note that both  $a$  and  $b$  are sub-gaussian random variables with sub-gaussian norm  $c/\sqrt{d}$  where  $c$  is some constant [Vershynin, 2010].

Denote  $\mathbb{E}_b[\cdot]$  the expectation over  $b$ . We can rewrite the probability as

$$\begin{aligned} \Pr\left[|\langle a, b \rangle| \geq \frac{cu}{\sqrt{d}}\right] &\leq \mathbb{E}_b \Pr\left[|\langle a, b \rangle| \geq \frac{cu}{\sqrt{d}} \mid b\right] \\ &\leq \mathbb{E}_b \{2 \exp(-u^2)\} = 2 \exp(-u^2). \end{aligned} \quad (49)$$

The last inequality uses the independence of  $a$  and  $b$  and  $\|\langle a, b \rangle\|_{\psi_2} \leq \|b\|_2 \|a\|_{\psi_2}$  for a fixed  $b$ .  $\square$

Decomposing the sum into diagonal and off-diagonal terms gives us

$$\|G_n - G\| \leq \sup_{\|\alpha\|=1} \sum_{i \neq j}^m \alpha_i \alpha_j \langle x_i, x_j \rangle E_{ij} + \sum_{i=1}^m \alpha_i^2 E_{ii} \quad (50)$$

$$\leq \sup_{\|\alpha\|=1} \sqrt{\sum_{i \neq j} \alpha_i^2 \alpha_j^2} \sqrt{\sum_{i \neq j} \langle x_i, x_j \rangle^2 E_{ij}^2} + \max_i |E_{ii}|. \quad (51)$$

Let  $\mathcal{G}$  denote the event that for all  $i \neq j \in [m]$ ,  $|\langle x_i, x_j \rangle| \leq O\left(\frac{\log d}{\sqrt{d}}\right)$ , then by Lemma 10 and the union bound

$$\Pr[-\mathcal{G}] \leq 2m^2 e^{-\log^2 d}. \quad (52)$$

Therefore, with probability at least  $1 - 2m^2 e^{-\log^2 d}$ , we have

$$\|G_n - G\| \leq c \frac{\log d}{\sqrt{d}} \sqrt{\sum_{i \neq j} E_{ij}^2} + \max_i |E_{ii}|. \quad (53)$$

Note that

$$U(\{x_1, \dots, x_m\}) = \frac{1}{m(m-1)} \sum_{i \neq j} E_{ij}^2 \quad (54)$$

is a U-statistics.

Suppose  $|E_{ij}| \leq B$ , according to the concentration inequality (Theorem 2 in [Peel et al., 2010]), we have with probability at least  $1 - \delta$

$$\sum_{i \neq j} E_{ij}^2 \leq m(m-1) \mathbb{E}_{\{x_1, x_2\}} E_{12}^2 + m(m-1) \left( \sqrt{\frac{4\Sigma^2}{m} \log \frac{1}{\delta}} + \frac{4B^2}{3m} \log \frac{1}{\delta} \right), \quad (55)$$

where  $\Sigma^2 = \mathbb{E}[E_{1,2}^4] - \mathbb{E}[E_{1,2}^2]^2$  is the variance for the kernel in U-statistics.

Putting everything together, we have with probability at least  $1 - \delta - 2m^2 e^{-\log^2 d}$

$$\|G_n - G\| \leq c \frac{\log d}{\sqrt{d}} \left( m \sqrt{\mathbb{E}_{\{x_1, x_2\}} E_{12}^2} + m \left( \frac{4\Sigma^2}{m} \log \frac{1}{\delta} \right)^{1/4} + mB \sqrt{\frac{4}{3m} \log \frac{1}{\delta}} \right) + B \quad (56)$$

## D Discrepancy of the weights

In this section, we relate the quantities  $\mathbb{E}_{\{x_1, x_2\}} E_{12}^2$  and  $B$  to the discrepancies of the weights. Note that for ReLU,  $\sigma'(w^\top x)$  does not depend on the norm of  $w$ , so we can focus on  $w$  on the unit sphere.

Given a set of  $n$  points  $W = \{w_i\}_{i=1}^n$  on the unit sphere  $\mathbb{S}^{d-1}$ , the discrepancy of  $W$  with respect to a measurable subset  $S \subseteq \mathbb{S}^{d-1}$  is defined as

$$\text{dsp}(W, S) = \frac{1}{n} |W \cap S| - A(S) \quad (57)$$

where  $A(S)$  is the normalized area of  $S$  (e.g., the area of the whole sphere  $A(\mathbb{S}^{d-1})$  is 1). Let  $\mathcal{S}$  denote the family of slices in  $\mathbb{S}^{d-1}$

$$\mathcal{S} = \{S_{xy} : x, y \in \mathbb{S}^{d-1}\}, \text{ where } S_{xy} = \{w \in \mathbb{S}^{d-1} : w^\top x \geq 0, w^\top y \geq 0\}. \quad (58)$$

The  $L_\infty$  discrepancy of  $W$  with respect to  $\mathcal{S}$  is

$$L_\infty(W, \mathcal{S}) = \sup_{S \in \mathcal{S}} \text{dsp}(W, S), \quad (59)$$

and the  $L_2$  discrepancy is

$$L_2(W, \mathcal{S}) = \sqrt{\mathbb{E}_{x, y} \text{dsp}(W, S_{xy})^2} \quad (60)$$

where the expectation is taken over  $x, y$  uniformly on the sphere. We use  $L_\infty(W)$  and  $L_2(W)$  as their shorthands.

For ReLU, by definition, we have

$$\mathbb{E} E_{ij}^2 = (L_2(W))^2, \quad (61)$$

$$B \leq L_\infty(W), \quad (62)$$

$$\Sigma^2 \leq \mathbb{E}[E_{1,2}^4] \leq \mathbb{E}[E_{1,2}^2] \max_{x_1, x_2} |E_{1,2}^2| \leq (L_\infty(W) L_2(W))^2, \quad (63)$$

using the fact that  $E_{ij} = \text{dsp}(W, S_{x_i x_j})$ .

Therefore, the bound becomes

$$\|G_n - G\| \leq c \frac{\log d}{\sqrt{d}} \left( mL_2(W) + \sqrt{L_\infty(W)L_2(W)}m \left( \frac{4}{m} \log \frac{1}{\delta} \right)^{1/4} + mL_\infty(W) \sqrt{\frac{4}{3m} \log \frac{1}{\delta}} \right) + L_\infty(W) \quad (64)$$

In the following subsections, we will discuss the discrepancies.

## D.1 Computing $L_2$ discrepancy for ReLU

Note that the derivative of ReLU  $\sigma'(w^\top x) = \mathbb{I}[w^\top x]$  does not depend on the norm of  $w$ . Without loss of generality, we can assume  $\|w\| = 1$  throughout this subsection.

**Theorem 7.** Suppose  $W = \{w_i\}_{i=1}^n \subseteq \mathbb{S}^{d-1}$ .

$$(L_2(W))^2 = \frac{1}{n^2} \sum_{i,j=1}^n k(w_i, w_j)^2 - \mathbb{E}_{u,v} [k(u, v)^2]$$

where  $\mathbb{E}_{u,v}$  is over  $u$  and  $v$  uniformly distributed on  $\mathbb{S}^{d-1}$  and the kernel  $k(\cdot, \cdot)$  is

$$k(u, v) = \frac{\pi - \arccos \langle u, v \rangle}{2\pi}.$$

*Proof.* Let  $d(u, v) = \frac{\arccos \langle u, v \rangle}{\pi}$ . Let  $S_{xy} = \{w \in \mathbb{S}^{d-1} : w^\top x \geq 0, w^\top y \geq 0\}$ . It is clear that (up to sets of measure zero)

$$A(S_{xy}) = k(x, y) = \frac{1 - d(x, y)}{2}, \quad (65)$$

$$\mathbb{I}[z \in S_{xy}] = \frac{1}{4} (\text{sign}(x^\top z) + 1) (\text{sign}(y^\top z) + 1), \quad (66)$$

where  $\mathbb{I}[\cdot]$  is the indicator function. Then

$$\text{dsp}(W, S_{xy}) = \frac{1}{n} \sum_{k=1}^n \mathbb{I}[w_k \in S_{xy}] - A(S_{xy}) \quad (67)$$

$$= \frac{1}{n} \sum_{k=1}^n \frac{1}{4} (\text{sign}(x^\top w_k) + 1) (\text{sign}(y^\top w_k) + 1) - \frac{1 - d(x, y)}{2}. \quad (68)$$

Let  $s_{xi}$  be a shorthand for  $\text{sign}(x^\top w_i)$ . Then we have

$$(L_2(W))^2 = \mathbb{E}_{x,y} (\text{dsp}(W, S_{xy}))^2 \quad (69)$$

$$= \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \left( \frac{1}{n} \sum_{k=1}^n \frac{1}{4} (\text{sign}(x^\top w_k) + 1) (\text{sign}(y^\top w_k) + 1) - \frac{1 - d(x, y)}{2} \right)^2 dA(x) dA(y) \quad (70)$$

$$= \frac{1}{n^2} \sum_{i,j=1}^n \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \frac{(s_{xi} + 1)(s_{yi} + 1)}{4} \frac{(s_{xj} + 1)(s_{yj} + 1)}{4} dA(x) dA(y) \quad (71)$$

$$- \frac{2}{n} \sum_{i=1}^n \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \frac{1 - d(x, y)}{2} \frac{(s_{xi} + 1)(s_{yi} + 1)}{4} dA(x) dA(y) \quad (72)$$

$$+ \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \left( \frac{1 - d(x, y)}{2} \right)^2 dA(x) dA(y). \quad (73)$$

Consider the first term, which is equal to

$$\frac{1}{n^2} \sum_{i,j=1}^n \left( \int_{\mathbb{S}^{d-1}} \frac{(s_{xi} + 1)(s_{xj} + 1)}{4} dA(x) \right) \left( \int_{\mathbb{S}^{d-1}} \frac{(s_{yi} + 1)(s_{yj} + 1)}{4} dA(y) \right). \quad (74)$$

By Lemma 11,

$$\int_{\mathbb{S}^{d-1}} \frac{(s_{xi} + 1)(s_{xj} + 1)}{4} dA(x) = \frac{2 - 2d(w_i, w_j)}{4}, \quad (75)$$

so the first term is equal to

$$\frac{1}{n^2} \sum_{i,j=1}^n \left( \int_{\mathbb{S}^{d-1}} \frac{(s_{xi} + 1)(s_{xj} + 1)}{4} dA(x) \right)^2 = \frac{1}{n^2} \sum_{i,j=1}^n k(w_i, w_j)^2. \quad (76)$$

Now consider the second term. Note that the summand is invariant to  $w_i$ , so it can be replaced by an arbitrary  $p \in \mathbb{S}^{d-1}$ . The second term is then equal to

$$-2 \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \frac{1 - d(x, y)}{2} \frac{(\text{sign}(x^\top p) + 1)(\text{sign}(y^\top p) + 1)}{4} dA(x) dA(y) \quad (77)$$

$$= -2 \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \frac{1 - d(x, y)}{2} \mathbb{I}[p \in S_{xy}] dA(x) dA(y) \quad (78)$$

$$= -2 \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \frac{1 - d(x, y)}{2} \mathbb{I}[p \in S_{xy}] dA(x) dA(y) dA(p) \quad (79)$$

$$= -2 \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \frac{1 - d(x, y)}{2} \left[ \int_{\mathbb{S}^{d-1}} \mathbb{I}[p \in S_{xy}] dA(p) \right] dA(x) dA(y) \quad (80)$$

$$= -2 \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \frac{1 - d(x, y)}{2} \frac{2 - 2d(x, y)}{4} dA(x) dA(y) \quad (81)$$

$$= -2 \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \left( \frac{1 - d(x, y)}{2} \right)^2 dA(x) dA(y), \quad (82)$$

where the third step is by invariance to  $p$  and the fourth step is by Lemma 11. The theorem then follows.  $\square$

Theorem 7 lets us compute  $L_2(W)$  for a fixed  $W$ . The next lemma gives a concrete bound for a special case where  $W$  is uniformly distributed on the unit sphere.

**Lemma 8** *There exists a constant  $c_g$ , such that for any  $0 < \delta < 1$ , with probability at least  $1 - \delta$  over  $W = \{w_i\}_{i=1}^n$  that are sampled from the unit sphere uniformly at random,*

$$(L_2(W))^2 \leq c_g \left( \sqrt{\frac{\log d}{nd} \log \frac{1}{\delta}} + \frac{1}{n} \log \frac{1}{\delta} \right).$$

*Proof.* By Theorem 7, we have

$$(L_2(W))^2 = \frac{1}{4n^2} \sum_{i,j=1}^n \left( \frac{1}{2} - d(w_i, w_j) \right)^2 - \frac{1}{4} \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \left( \frac{1}{2} - d(u, v) \right)^2 dA(u) dA(v) + \frac{1}{4n^2} \sum_{i,j=1}^n \left( \frac{1}{2} - d(w_i, w_j) \right). \quad (83)$$

First consider  $T_1 = \frac{1}{n^2} \sum_{i,j=1}^n \left( \frac{1}{2} - d(w_i, w_j) \right)^2 - \mu$  where  $\mu = \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \left( \frac{1}{2} - d(u, v) \right)^2 dA(u) dA(v)$ . Rewrite  $T_1 = \frac{1}{4n} + \frac{n-1}{n} U(W) - \mu$  where  $U(W) = \frac{1}{n(n-1)} \sum_{i \neq j} \left( \frac{1}{2} - d(w_i, w_j) \right)^2$  is a U-statistics. We upper bound  $U(W)$  by using Bernstein's inequality when  $W$  is uniform over the sphere.

By Taylor expansion, we have

$$\frac{1}{2} - d(u, v) = x/\pi + x^3/6\pi + O(x^5), \text{ where } x = u^\top v.$$

Then let  $\mathcal{G}$  denote the event that  $|x| = |u^\top v| \leq c\sqrt{\log d/d}$  for a sufficient large constant  $c > 0$ , so that by Lemma 10,

$\Pr[\neg\mathcal{G}] \leq O(1/d^4)$ . Then

$$\mathbb{E}[U(W)] = \mu = \mathbb{E}\left[\left(x/\pi + x^3/6\pi + O(x^5)\right)^2\right] \quad (84)$$

$$= \mathbb{E}[x^2/\pi^2 + x^4/6\pi^2 + O(x^6)] \quad (85)$$

$$\leq \mathbb{E}\left[\left[x^2/\pi^2 + x^4/6\pi^2 + O(x^6)\right] \mid \mathcal{G}\right] + \Pr[\neg\mathcal{G}] \max_{u,v} \left[\frac{1}{2} - d(u,v)\right]^2 \quad (86)$$

$$= O\left(\frac{\log d}{d}\right), \quad (87)$$

and thus

$$\text{Var}[U(W)] = \mathbb{E}\left\{\left[\left(x/\pi + x^3/6\pi + O(x^5)\right)^2 - \mu\right]^2\right\} \quad (88)$$

$$= \mathbb{E}\left\{\left[x^2/\pi^2 + x^4/6\pi^2 + O(x^6) - \mu\right]^2\right\} \quad (89)$$

$$\leq \mathbb{E}\left\{\left[x^2/\pi^2 + x^4/6\pi^2 + O(x^6) - \mu\right]^2 \mid \mathcal{G}\right\} + \Pr[\neg\mathcal{G}] \max_{u,v} \left[\left(\frac{1}{2} - d(u,v)\right)^2 - \mu\right]^2 \quad (90)$$

$$= O\left(\frac{\log^2 d}{d^2}\right). \quad (91)$$

Then by Bernstein's inequality, we have with probability at least  $1 - \delta$  over the  $W$  uniformly on the sphere,

$$|T_1| \leq O\left(\frac{\log d}{d} \sqrt{\frac{1}{n} \log \frac{1}{\delta}} + \frac{1}{n} \log \frac{1}{\delta}\right). \quad (92)$$

A similar argument holds for  $T_2 = \frac{1}{n^2} \sum_{i,j=1}^n \left(\frac{1}{2} - d(w_i, w_j)\right)$ . Note that

$$\text{Var}\left\{\left(\frac{1}{2} - d(u,v)\right)\right\} = \mu = O\left(\frac{\log d}{d}\right). \quad (93)$$

We have that with probability at least  $1 - \delta$  over the  $W$  uniform from the sphere,

$$|T_2| \leq O\left(\sqrt{\frac{\log d}{nd} \log \frac{1}{\delta}} + \frac{1}{n} \log \frac{1}{\delta}\right). \quad (94)$$

This completes the proof.  $\square$

Below are some technical lemmas that are used in the analysis.

**Lemma 11.**

$$\int_{\mathbb{S}^{d-1}} d(x,y) dA(x) = \frac{1}{2}, \forall y \in \mathbb{S}^{d-1}, \quad (95)$$

$$\int_{\mathbb{S}^{d-1}} \text{sign}(x^\top y) dA(x) = 0, \forall y \in \mathbb{S}^{d-1}, \quad (96)$$

$$\int_{\mathbb{S}^{d-1}} \text{sign}(x^\top z) \text{sign}(y^\top z) dA(z) = 1 - 2d(x,y), \forall x,y \in \mathbb{S}^{d-1}. \quad (97)$$

*Proof.* The first two are straightforward. The third is implicit in the proof of Theorem 1.21 in [Bilyk and Lacey, 2015].  $\square$

## E The spectrum of $\gamma_m$

The spectrum of the kernel matrix  $g(x,y) = \left(\frac{1}{2} - \frac{\arccos\langle x,y \rangle}{2\pi}\right) \langle x,y \rangle$  is determined by the spherical decomposition coefficients.

We need  $\gamma_m$  to decrease slower than  $O(1/\sqrt{m})$  within a reasonable range, such as  $m \leq 1000000$ .

Although the kernel associated with ReLU decreases faster than the desired rate, we can choose from a family of such arccos kernels such that the spectrum decays slower than  $1/\sqrt{m}$ .

Such arccos kernels are defined as

$$k_n(x, y) = J_n(\theta)/\pi \quad (98)$$

where

$$J_n(\theta) = (-1)^n (\sin \theta)^{2n+1} \left( \frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \right)^n \left( \frac{\pi - \theta}{\sin \theta} \right) \quad (99)$$

Larger  $n$  corresponds to more nonlinear activation functions and leads to slower decaying spectrum.

$$\begin{aligned} J_0(\theta) &= \pi - \theta \\ J_1(\theta) &= \sin \theta + (\pi - \theta) \cos \theta \\ J_2(\theta) &= 3 \sin \theta \cos \theta + (\pi - \theta)(1 + 2 \cos^2 \theta) \\ 4J_3(\theta) &= 27 \sin \theta + 11(3 \sin \theta \cos^2 \theta - \sin^3 \theta) \\ &\quad + (\pi - \theta)(54 \cos \theta + 6(\cos^3 \theta - 3 \sin^2 \theta \cos \theta)) \\ 4J_4(\theta) &= (\pi - \theta)(216 + 192 \cos(2\theta) + 12 \cos(4\theta)) \\ &\quad + 160 \sin(2\theta) + 25 \sin(4\theta) \\ 8J_5(\theta) &= (\pi - \theta)(6000 \cos \theta + 1500 \cos(3\theta) + 60 \cos(5\theta)) \\ &\quad + 2000 \sin \theta + 1625 \sin(3\theta) + 137 \sin(5\theta) \\ \frac{8}{9}J_6(\theta) &= (\pi - \theta)(4000 + 4500 \cos(2\theta) + 720 \cos(4\theta) + 20 \cos(6\theta)) \\ &\quad + 2625 \sin(2\theta) + 924 \sin(4\theta) + 49 \sin(6\theta) \\ \frac{16}{9}J_7(\theta) &= (\pi - \theta)(171500 \cos(\theta) + 61740 \cos(3\theta) + 6860 \cos(5\theta) + 140 \cos(7\theta)) \\ &\quad + 42875 \sin(\theta) + 48363 \sin(3\theta) + 9947 \sin(5\theta) + 363 \sin(7\theta) \end{aligned} \quad (100)$$

Higher orders of  $J_n(\theta)$  seems to be extremely complicated.

Although there is no analytical solution to the spectrum, we can compute them numerically.

Figure 1 illustrates the spectra of several arccos kernels compared to  $O(1/m)$  and  $O(1/\sqrt{m})$ .

## F Rademacher complexity and final error bounds: Proof of Theorem 2 and Theorem 1

We apply the argument in [Bartlett and Mendelson, 2002] to our setting to get Lemma 12. Combining it with Theorem 3 leads to Theorem 2. Further combining it with Lemma 8 leads to Theorem 1.

**Lemma 12.** *Suppose the data are bounded:  $|y| \leq Y$  and  $\|x\|_2 \leq 1$ . Let*

$$\mathcal{F} = \left\{ f(x) = \sum_{k=1}^n v_k \sigma(w_k^\top x) : v_k \in \{-1, +1\}, \sum_k \|w_k\|_2 \leq C_W \right\}.$$

*Then with probability  $\geq 1 - \delta$ , for any  $f \in \mathcal{F}$ ,*

$$\frac{1}{2} \mathbb{E}(y - f(x))^2 \leq \frac{1}{2m} \sum_{l=1}^m (y_l - f(x_l))^2 + \frac{2(Y + C_W)C_W}{\sqrt{m}} + (Y^2 + C_W^2) \sqrt{\frac{\log \frac{1}{\delta}}{2m}}. \quad (101)$$

*Proof.* For a sample  $S = ((x_1, y_1), \dots, (x_m, y_m))$ , and a loss function  $\ell(y, x) = \frac{1}{2}(y - f(x))^2$ , we denote  $\hat{\mathbb{E}}_S[\ell]$  as the empirical average  $\hat{\mathbb{E}}_S[\ell] = \frac{1}{m} \sum_{l=1}^m \ell(y_l, x_l)$ .

Define

$$\Phi(S) = \sup_{\ell \in \mathcal{L}} \mathbb{E}[\ell] - \hat{\mathbb{E}}_S[\ell] \quad (102)$$

where  $\mathcal{L}$  is the set of loss functions

$$\mathcal{L} = \left\{ \ell(y, x) = \frac{1}{2}(y - f(x))^2 : f \in \mathcal{F} \right\}.$$

Let  $S$  and  $S'$  be two datasets that differ by exactly one data point  $(x_i, y_i)$  and  $(x'_i, y'_i)$ . Then we have a bound on the difference of loss functions. Since  $\|x\|_2 \leq 1$  and  $\sum_k \|w_k\|_2 \leq C_W$ , we have  $|f| \leq C_W$ . Thus

$$|\ell(y, f(x)) - \ell(y', f(x'))| \leq \frac{1}{2} \max \{ (y - f(x))^2, (y' - f(x'))^2 \} \leq Y^2 + C_W^2. \quad (103)$$

This leads to an upper bound

$$\begin{aligned} \Phi(S) - \Phi(S') &\leq \sup_{\ell \in \mathcal{L}} \hat{\mathbb{E}}_S[\ell] - \hat{\mathbb{E}}_{S'}[\ell] \\ &= \sup_{\ell \in \mathcal{L}} \frac{\ell(y_i, f(x_i)) - \ell(y'_i, f(x'_i))}{m} \leq \frac{Y^2 + C_W^2}{m}. \end{aligned} \quad (104)$$

Similarly, we can get the other side of the inequality and have  $|\Phi(S) - \Phi(S')| \leq \frac{Y^2 + C_W^2}{m}$ . From McDiarmids' inequality, with probability at least  $1 - \delta$  we get

$$\Phi(S) \leq \mathbb{E}_S \Phi(S) + (Y^2 + C_W^2) \sqrt{\frac{\log \frac{1}{\delta}}{2m}}. \quad (105)$$

The first term on the right-hand side can be bounded by Rademacher complexity as shown in the book Foundations of Machine Learning (3.13). In the end, we have the bound

$$\frac{1}{2} \mathbb{E}(y - f(x))^2 \leq \frac{1}{2m} \sum_{l=1}^m (y_l - f(x_l))^2 + 2\mathcal{R}_m(\mathcal{L}) + (Y^2 + C_W^2) \sqrt{\frac{\log \frac{1}{\delta}}{2m}} \quad (106)$$

where  $\mathcal{R}_m(\mathcal{L})$  is the Rademacher complexity of the function class  $\mathcal{L}$ .

We can find the Rademacher complexity by using composition rules. The Rademacher complexity of linear functions  $\{w^\top x : \|w\|_2 \leq b_W, \|x\|_2 \leq 1\}$  is  $b_W / \sqrt{m}$ , where  $m$  is the number of data points. If a function  $\phi$  is  $L$ -Lipschitz, then for any function class  $\mathcal{H}$ , we have  $\mathcal{R}(\phi \circ \mathcal{H}) \leq L\mathcal{R}(\mathcal{H})$ . In addition, we also have  $\mathcal{R}(c\mathcal{H}) = |c|\mathcal{R}(\mathcal{H})$  and  $\mathcal{R}(\sum_k F_k) \leq \sum_k \mathcal{R}(F_k)$ .

So for the function class  $\mathcal{F}$  that describes a neural network, we have

$$\mathcal{R}_m(\mathcal{F}) \leq \frac{C_W}{\sqrt{m}}. \quad (107)$$

It is derived by using the fact that  $\sigma'(\cdot)$  is 1-Lipschitz and  $\sum_k \|w_k\|_2 \leq C_W$ .

Finally composing on the loss function we get

$$\mathcal{R}_m(\mathcal{L}) \leq \frac{(Y + C_W)C_W}{\sqrt{m}}, \quad (108)$$

using the fact that the ground truth in the loss should be bounded by  $Y$  and the function bounded by  $C_W$ , thus the Lipschitz constant of the loss function is bounded by  $Y + C_W$ .  $\square$